

calibrar: an R package for fitting complex ecological models

Ricardo Oliveros-Ramos^{1,2,3} and Yunne-Jai Shin^{2,3}

¹Instituto del Mar del Perú (IMARPE). Gamarra y General Valle s/n Chucuito, Callao, Perú.

²Institut de Recherche pour le Développement (IRD), UMR MARBEC 248, Université de Montpellier, place E Bataillon, Bât 24, CC093, 34095 Montpellier cedex 5, France.

³University of Cape Town, Marine Research (MA-RE) Institute, Department of Biological Sciences, Private Bag X3, Rondebosch 7701, South Africa.

E-mail Ricardo Oliveros-Ramos: ricardo.oliveros@gmail.com

Summary

1. The fitting or parameter estimation of complex ecological models is a challenging optimization task, with a notable lack of tools for fitting complex stochastic models.
2. calibrar is an R package that has been developed for fitting complex ecological models to data, including complex Individual Based Models. It is a generic tool that can be used for any type of model, especially those with non-differentiable objective functions.
3. calibrar supports multiple phases and constrained optimization. It implements maximum likelihood estimation methods and automated construction of the objective function from simulated model outputs.
4. User-level expertise in R is necessary to handle calibration experiments with calibrar, but there is no need to modify the model's code, which can be programmed in any language. For more experienced users, calibrar allows the implementation of user-defined objective functions.
5. The package source code is fully accessible and can be installed directly from CRAN.

Keywords: black-box optimization, inverse problem, parameter estimation, calibration, evolutionary algorithms, stochastic model, individual based model, OSMOSE model.

Introduction

The ability to achieve accurate parameter estimation has been used as a criterion to assess the usefulness of ecological models (Bartell 2003). Given a model, the criterion for the selection of the best possible parameter set is the optimization of a scalar *objective function* (e.g. log-likelihood, residual sum of squares) with respect to the model parameters (Walter and Pronzato 1997, Bolker et al. 2013). Once the objective function is properly defined, parameter estimation is essentially an optimization problem. The parameter estimation or calibration of complex ecological models (Jorgensen XXX) could be a difficult task for optimization because of model characteristics such as non-linearity, high dimensionality as well as low quantity and quality of observed data (Tashkova et al. 2012). These diverse factors have hampered the development of calibration algorithms and methodologies for ecological models that are sufficiently flexible and generic, and only sparse documentation has been produced on fitting complex models (Bolker et al. 2013). Additionally, complex ecosystem models can be numerically intensive and require long simulation runs, adding an extra layer of difficulty to their fitting.

There are some dedicated tools for non-linear parameter estimation, AD Model Builder (ADMB; Fournier et al. 2012) being one of the most robust and fast (Bolker et al. 2013). Among other advantages, ADMB provides support for parameter estimation in multiple phases (Nash and Walker-Smith 1987), which can be of great interest when dealing with complex ecosystem models (Oliveros-Ramos et al. 2015). It also provides support for constraining optimization, which can be helpful for regularizing hard optimization problems (Bolker et al. 2013). However, the model and the objective function itself need to be coded in C++ (using the ADMB scripting), which can be an obstacle for fitting complex models which have been already implemented in other languages (e.g. Java, Fortran). In addition, as ADMB is based on automatic differentiation (AD), which allows to provide accurate estimates of derivatives (Griewank and Corliss 1992), but does not handle the estimation of parameters of stochastic models for which derivatives cannot be computed.

Parameter estimation methods have been developed for stochastic non-linear models, such as continuous time, finite state Markov models and individual-based models (IBMs), for which the probability of state transitions or the master equation can be written (Ionides et al. 2006, Newman et al. 2009, Ross et al. 2009, Walker et al. 2006). However, many stochastic models at the individual level can only be simulated numerically and are too complex for mathematical analysis and explicit parameter estimation (Black and McKane 2012), resulting in more attention being given to the exploration of model behaviour than to a rigorous confrontation with data. To solve these issues, meta-heuristic algorithms (e.g. evolutionary algorithms) have been used (Cropper and Anderson 2004, Poovathingal and Gunawan 2010, Duboz et al. 2010, Tashkova et al. 2012, Travers-Trolet et al. 2013), and have in some cases shown better performance than derivative-based optimization methods (Tashkova et al. 2012). However, the scientific community lacks generic, open and flexible enough tools for the parameter estimation of different types of ecological models with different degrees of complexity.

Here we present a novel R package, *calibrar*, designed for the parameter estimation of complex ecological models, including stochastic ones. The package provides support for multiple phases and parameter constraint optimization. In particular, by using a “black-box” approach, the package allows the calibration of models implemented in any programming language. It provides a generic interface with models and allows the construction of the objective function, within R, without requiring any changes in the models' code. Parallel support for computationally intensive models is also provided, and can be used with high performance computing systems.

General description of the package

The `calibrar` package is written in R (R Development Core Team 2014), and can be installed from CRAN (`install.packages("calibrar")`). The package was primarily designed to handle the parameter estimation of complex models using a maximum likelihood approach. However, it can be used for fitting any model, and can also be used with user-defined objective functions (e.g. least squares). The package was designed for the optimization of “black-box” functions (Jones et al. 1998), where analytical information about the function to be optimized and the model source code are assumed to be unavailable or impractical to modify (Rios and Sahinidis 2013). Our approach is hence “non-intrusive”, making the model interact with the optimizer, i.e. the `calibrar` package, in two ways: i) receiving a set of parameters to run the model, and ii) providing the model outputs to be confronted with the observed data. The package also helps in the construction of the objective function to be optimized in order to estimate model's parameters (Figure 1).

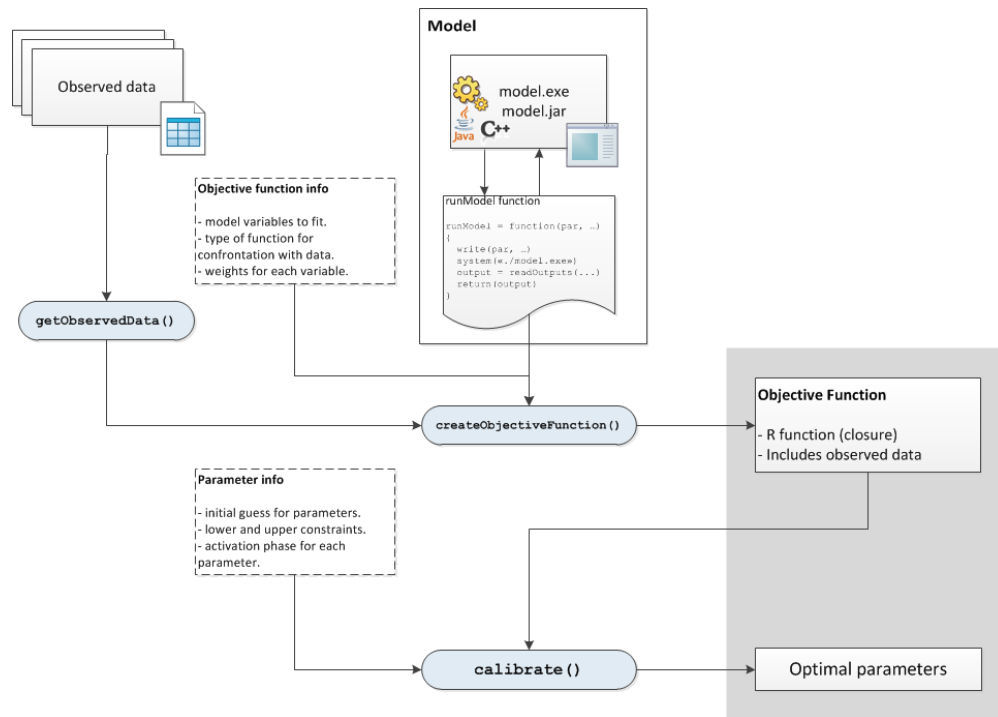


Figure 1. Diagram representing the functioning of the `calibrar` package. The grey area groups the outputs produced by the package (the objective function and the optimal parameters of the model). Rectangles with broken border lines show user inputs which are needed to configure the optimization. Rounded rectangles show main package functions.

The package was designed in a way that minimal expertise in R is necessary to handle the model fitting and the main functionalities are embedded in three functions: `calibrate`, `getObservedData` and `createObjectiveFunction` (see Table 1). The user intervention is mainly required in the construction of the function to run the model (`runModel` function, Figure 1) and in retrieving the simulation outputs within R. However, given R’s flexibility and features for data manipulation, it is rather straightforward to develop such a function. Some complex ecological models have dedicated packages oriented to the analysis and simulation of their outputs that could be used to link with `calibrar`, e.g., `RNetLogo` (Thiel et al. 2012) for IBMs implemented in `netLogo` or `osmose2R` (www.osmose-model.org) for the OSMOSE model (Shin and Cury 2001, 2004). The main function of the package is `calibrate`, which performs minimization of the objective function. It has a similar syntax as the classical R optimization function `optim` (see Table 2), and is also similar to most optimization related functions in other R packages. Additionally, the `getObservedData` and

`createObjectiveFunction` are provided to simplify the organization of the observed data and construct the objective function for the calibration, respectively.

Table 1. Functions of the `calibrar` package.

Function	Returned objects	Description
<code>calibrate</code>	An object summarizing the calibration results	Performs a sequential parameter estimation of a model using multiple phases
<code>getObservedData</code>	A list with the observed values (the data)	Create a list with the observed data with the information provided by its main argument
<code>createObjectiveFunction</code>	A function, integrating the simulation of the model and the comparison with observed data	Create a new function, to be used as the objective function in the calibration, given a function to run the model within R, observed data and information about the comparison with data
<code>getCalibrationInfo</code>	A <code>data.frame</code> with the information necessary to create the objective function using <code>createObjectiveFunction</code>	Basically a wrapper for <code>read.csv</code> checking column names and data types for the table with the calibration information
<code>coef</code> , <code>summary</code> , <code>predict</code>		R S3 methods for visualizing the results of the calibration

In order to create the objective function, the user needs to specify some information about the model outputs used for the calibration and how to combine them (Figure 1). More experienced users can create the objective function by directly integrating the simulation of the model (the main purpose of the `runModel` function). The details for the creation of the objective function are explained in the next section.

Additionally, the user needs to specify information on the parameters to calibrate. We recommend that lower and upper thresholds are provided for each parameter, but unconstrained optimization is also supported. If initial approximate estimates for the parameter values are provided, this would simplify the calibration process, but this step can be omitted if no information is available. In case of a multiple phase calibration, the user must indicate the phase of the calibration where the estimation of a parameter must be included. The implementation of multiple phases in the calibration is detailed in the next section.

The `calibrar` package allows to perform the calibration with a total of 18 different optimization algorithms, including the ones available in the package `stats` (R Core Team 2016), `optimx` (Nash and Varadhan 2011) and `cmaes` (Trautmann et al. 2011). The default option is AHR-ES which designates a novel algorithm based on Evolutionary Strategies (ES) that we specifically developed for the global optimization of complex stochastic models (described in appendix 1).

Table 2. Main arguments of the `calibrate` function.

Argument	Description
<code>par</code>	A numeric vector or a list containing numeric vectors. The length of the <code>par</code> argument defines the number of parameters to be estimated (i.e. the dimension of the problem)
<code>fn</code>	The function to be minimized
<code>gr</code>	A function specifying the gradient of <code>fn</code>
<code>upper</code>	Upper threshold value(s) for parameters. One value or a vector of the same length as <code>par</code> . If one value is provided, it is used for all parameters. NA means <code>Inf</code> . By default <code>Inf</code> is used (unconstrained).
<code>lower</code>	Lower threshold value(s) for parameters. One value or a vector of the same length as <code>par</code> . If one value is provided, it is used for all parameters. NA means <code>-Inf</code> . By default <code>-Inf</code> is used (unconstrained).
<code>phases</code>	An optional vector of the same length as <code>par</code> , indicating the phase at which each

	parameter becomes active. If omitted, default value is 1 for all parameters. Negative integers and NA are accepted for phases, both meaning that the parameter will never be active, so it will remain constant throughout the calibration.
<code>replicates</code>	The number of replicates of model simulation to evaluate the objective function in case of stochastic models. One value or a vector of length <code>max(phases)</code> , to specify a different number of replicates for each phase. The default value is 1.
<code>method</code>	The optimization method to be used, 18 methods are currently available
<code>aggFn</code>	Default is <code>weighted.sum</code> (and default weights all equal to 1).
<code>control</code>	Fine control of the optimization, see function help for details and Table 3.

Linking to models and data

The main purpose of the package is to fit complex models to data. In order to solve a calibration problem, we first need to define the objective function for our problem. As mentioned before, a non-intrusive black-box optimization approach is adopted, which means that the computer code of the model to calibrate does not require to be modified, but the model can be evaluated for a given set of parameters. For this purpose, we developed an R function to i) write a set of parameters in the format the model is able to read, ii) run the model with this set of parameters and iii) read the model outputs back into R (Figure 2). The output of this `runModel` function is expected to be a list, each element being one of the variables to calibrate. If the model is already implemented in R, the construction of this function should be very simple. On the other hand, R facilities to process and analyse data in different formats allow to handle model outputs independently of the language used for coding the model.

After the construction of the `runModel` function, the second step consists in providing information for the construction of the objective function. Each variable listed in output of `runModel` needs to be documented in the objective function (name of "variable", "type", "calibrate", "weight" and "useData"; Figure 3). This information should be provided as a `data.frame`, and will be used as an argument for the functions `getObservedData` and `createObjectiveFunction`. The `getObservedData` function is expected to read data from the disk, to produce a list with the same structure as the outputs of the `runModel` function. The function `createObjectiveFunction` will combine the observed data and the `runModel` function to create the objective function for the calibration problem (Figure 3), which in turn will be the `fn` argument for the `calibrate` function.




Figure 2. Scheme of the link between the model and the calibration. The R function `runModel` receives a vector of parameters to test, writes the parameters in a form readable for the model (e.g. via txt or csv files), runs the model (possibly via `system`) then captures and processes the model outputs. The result of the function is a “list” object with all the variables to be confronted to observed data.

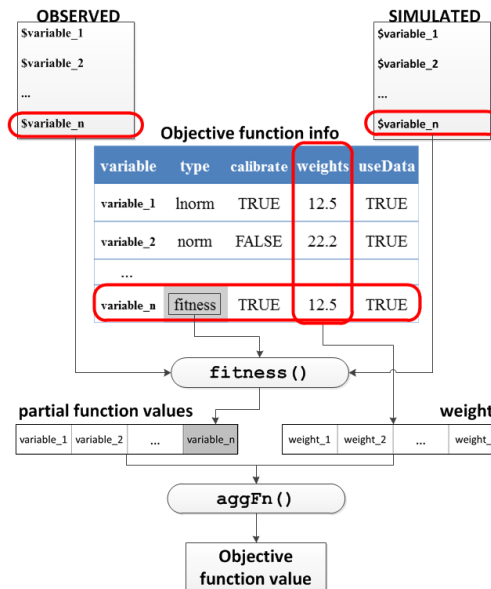


Figure 3. Scheme of the calculation of the value of the objective function for a given set of parameters. For each variable, a partial value of the objective function is calculated by applying the function specified in the column ‘type’ to the observed and simulated values. The final value of the objective function is calculated by applying the `aggFn` to the partial function values and the weights specified in the “objective function info” table.

To build the objective function, the ‘type’ selected for each variable is the function that will combine the observed and simulated data to produce a scalar value, measuring the fit between the model and the

observations. Some negative log-likelihood functions are already implemented and proposed for common distributions (e.g. normal, lognormal, multinomial, poisson; type `?fitness` to see the available functions). User defined functions can be provided, as they accept two vector arguments (`obs` and `sim`) and return a scalar value. For example, to implement a calibration using the least squares method, we can write the following function to calculate the residual sum of squares:

```
RSS = function(obs, sim, ...) {
  value = sum((obs-sim)^2, na.rm=TRUE)
  return(value)
}
```

The ‘calibrate’ column in the objective function information provides flags to select the variables to be used for the calibration. The ‘useData’ column indicates whether data are read from the disk. If `useData=TRUE`, a file called `variable_n.csv` will be searched. If `useData=FALSE`, the observed value is set to `NULL`, and the type function is expected to use simulated data only. The latter option can be particularly useful to set penalties in the model outputs or parameters, where no observed data are needed. Finally, the ‘weight’ column provides the relative weights to combine the partial objective values obtained for each variable. A more detailed illustration of this process is provided in appendix S2.

The `calibrate` function takes a list as a control argument, where fine control options are provided, for example the parallelization of the optimization. This feature is implemented for the default optimization method AHR-ES and is based on the `foreach` package (Revolution Analytics and Weston 2014). Before using the parallel implementation, a parallel ‘cluster’ should be created, which can be easily done using the `parallel` or `snow` R packages (see appendix S2 for an example). They allow the user to fully control the configuration of the parallel runs, making the calibration work in different computer systems, from computers with multicore processors to high-performance supercomputers. Once the cluster is created, only the `parallel=TRUE` and `ncores` control options should be provided (see Table 3 for details). Additionally, since each model run could require files to be written to the disk (which will be read by the `runModel` function after the simulation), a different folder needs to be assigned for each parameter combination that is tested by the optimization algorithm. For this purpose, the `run` control option allows a directory to be specified where all the simulations are run (subfolders named `i0`, `i1`, ..., `in` will be automatically created for `n` combinations tested in parallel). By default, no folders are created, so a path should be specified if the model needs to write files to the disk. All the parameter input files (Figure 2) will be written in temporary folders (e.g. `run/i0`). The control option `master` allows to specify a folder, the full content of which will be copied to temporary folders. Only files that need to be changed between individual runs are recommended to be put in the master folder; the use of absolute or relative paths is recommended for common heavy files needed to run the model. Since the calibration of numerically intensive models can run for a long time, a ‘restart’ option is also available, allowing an interrupted calibration to be continued.

Table 3. Some options for the `control` argument of the function `calibrate`.

Option	Description
<code>maxit</code>	Maximum number of executions of the objective function
<code>maxgen</code>	Maximum number of generations for the AHR-ES optimization algorithm. Ignored if <code>maxit</code> is provided, and recalculated accordingly.
<code>parallel</code>	Boolean, <code>TRUE</code> or <code>FALSE</code> to activate the parallel execution of the optimization.
<code>ncores</code>	The number of cores available in the parallel cluster for the active session. If <code>parallel=TRUE</code> , the default is to get the number of cores of the system.
<code>run</code>	An optional folder path to create all the temporary folders needed to run the simulations for each parameter combination tested by the optimization algorithm. The folders are recycled every generation.
<code>master</code>	An optional folder path. All the contents of the designated folder will be copied

	to each temporary folder.
save	Number of generations after saving a new restart object, which contains all the information necessary to restart the calibration at that point. The default is NULL, and no restart files are created.
restart.file	Filename for the restart file to be created.

Three application examples

To illustrate the main functionality of the package, we estimated the parameters for a predator-prey model using the `calibrate` function. The model was defined by a system of ordinary differential equations for the abundance of prey N and predator P :

$$\begin{aligned}\frac{dN}{dt} &= rN \left(1 - \frac{N}{K}\right) - \alpha NP \\ \frac{dP}{dt} &= -lP + \gamma \alpha NP\end{aligned}$$

The parameters to estimate were the prey's growth rate r , the predator's mortality rate l , the carrying capacity of the prey K and α and γ for the predation interaction. To start, we created the demonstration data for this model using the function `calibrarDemo` with $T=100$ as an additional argument to specify the time horizon.

```
demo          = calibrarDemo(path=path, model="PredatorPrey", T=100)
calibrationInfo = getCalibrationInfo(path=demo$path)
observed       = getObservedData(info=calibrationInfo, path=demo$path)
runModel       = calibrar:::PredatorPreyModel
obj            = createObjectiveFunction(runModel=runModel, info=calibrationInfo,
                                         observed=observed, T=demo$T)
```

To run the calibration, we needed to specify the initial guess for the parameter values (`par`), the objective function to minimize (`fn`) and optionally the lower and upper thresholds for the parameters (`lower` and `upper`) and the phase number at which each parameter needs to be estimated (`phases`).

```
calibrate(par=demo$guess, fn=obj, lower=demo$lower, upper=demo$upper, phases=demo$phase)
```

The argument `method` can be used to change the default optimization algorithm. We calibrated the model using five different optimization algorithms: AHR-ES (default, see appendix S1), the conjugate gradient (CG), L-BFGS-B and Nelder-Mead (default method in `optim`) as implemented in `stats::optim`, and the CMA-ES (implemented in the `cmaes` package, Trautmann et al. 2011). The results are shown in Figure 4 and Table 4, where we can observe that the algorithms CG and Nelder-Mead could not find the solution using the default parameters provided in the original optimization functions (while some improvements may be obtained after some tuning). The full code can be found in the appendix S2, section 4.




Figure 4. Results of the calibration of the predator-prey model using different optimization methods. The simulated data (points) and model fits (lines) are shown. For the methods “AHR-ES”, “L-BFGS-B” and “CMA-ES” there are no visual differences and the lines merge.

Table 4. Summary of the calibration results for the predator-prey model using five different optimization methods with default parameters.

Method	Objective function value	Parameter				
		r	l	K	α	γ
data	5.5661E-07	0.5000	0.2000	100.000	0.1000	0.1000
AHR-ES	5.7463E-04	0.4966	0.2020	99.5521	0.0992	0.1019
CG	2.5479E+00	0.4474	0.3173	56.2927	0.0734	0.2285
L-BFGS-B	1.3411E-05	0.4995	0.2003	99.9259	0.0999	0.1003
CMA-ES	6.5128E-07	0.5000	0.2000	99.9985	0.1000	0.1000
Nelder-Mead	4.7321E+01	0.5370	0.1241	59.6254	0.0494	0.1019

An additional example involves the calibration of a Poisson Autoregressive Mixed model for the dynamics of a population in different sites:

$$\log(\mu_{i,t+1}) = \log(\mu_{i,t}) + \alpha + \beta X_{i,t} + \gamma_t,$$

where $\mu_{i,t}$ is the size of the population in site i at year t , $X_{i,t}$ is the value of an environmental variable in site i at year t . The parameters to estimate were α , β , and γ_t , the random effects for each year ($\gamma_t \sim N(0, \sigma^2)$), and the initial population at each site $\mu_{i,0}$. We assumed that the observations $N_{i,t}$ follow a Poisson distribution with mean $\mu_{i,t}$. We could also create the data for this model using the function `calibrarDemo`, with the additional arguments `L=5` (five sites) and `T=100` (one hundred years):

```
demo = calibrarDemo(path=path, model="PoissonMixedModel", L=5, T=100)
calibrationInfo = getCalibrationInfo(path=demo$path)
```

```
observed      = getObservedData(info=calibrationInfo, path=demo$path)
forcing       = read.csv(file.path(demo$path, "master", "environment.csv"), row.names=1)
```

For this example, the `runModel` function was a bit more complicated, since the time series of γ_t was needed to construct the objective function:

```
runModel = function(par, forcing) {
  output = calibrar:::PoissonMixedModel(par=par, forcing=forcing)
  output = c(output, list(gammas=par$gamma))
  return(output)
}

obj = createObjectiveFunction(runModel=runModel, info=calibrationInfo, observed=observed,
forcing=forcing)
```

For this example we also specified different weights for each variable and increased the maximum number of iterations since a total of 107 parameters were estimated so the optimization was more complex:

```
control = list(weights=calibrationInfo$weights, maxit=3.6e5)
```

The calibration was run as in the previous example calling the `calibrate` function:

```
calibrate(par=demo$guess, fn=obj, lower=demo$lower, upper=demo$upper, control=control)
```

Here we calibrated the model using five different optimization algorithms: AHR-ES, CG, L-BFGS-B and SANN (Simulated Annealing from `optim`) and the CMA-ES. The results are shown in Figure 5 and Table 5. In this case, the SANN and CMA-ES could not find the solution despite the increased number of iterations. The best solution found was using the AHR-ES algorithm, and although the α and β parameters estimates were more far from the real values as the ones obtained with the CG and L-BFGS-B algorithms, the estimation of the distribution of the γ_t was closer to the real distribution and less biased (Figure 5b). The full code can be found in the appendix S2, section 4.

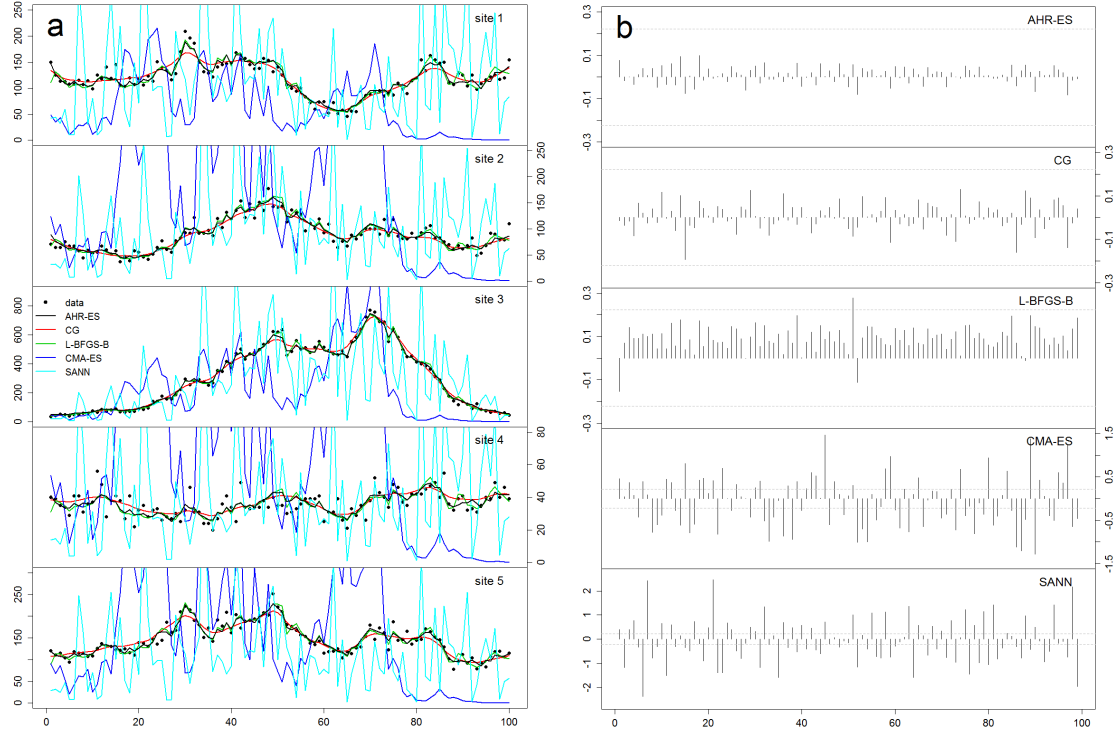


Figure 5. Results of the calibration of the autoregressive Poisson mixed model using different optimization methods. (a) The simulated data (points) and model fits (lines) are shown. (b) Time series of differences between the γ_t estimated by each algorithm and the real parameter values. For each case, the dotted lines represent the 95% limits for the differences assuming the original distribution.

Table 5. Summary of the calibration results for the autoregressive Poisson mixed model using five different optimization methods.

Method	Objective function value	Parameters					
		α	β	γ_t			
				mean	sd	cor	bias
data	-299310.600	0.4000	-0.4000	0	0.08	1.0000	0.0000
AHR-ES	-299306.795	0.3684	-0.3684	-0.0255	0.0768	0.7664	0.0285
CG	-299225.057	0.3902	-0.3896	0.0000	0.0293	0.6330	0.0030
L-BFGS-B	-299218.088	0.3824	-0.3807	-0.0885	0.0951	0.7792	0.0915
CMA-ES	-243318.074	0.1706	-0.2299	-0.0986	0.5322	0.0479	0.1016
SANN	-253430.498	0.3225	-0.3154	-0.4207	1.4749	0.0003	0.4237

A more complex application involved the calibration of the OSMOSE model (Shin and Cury 2001, 2004) in the Northern Humboldt Current Ecosystem and the Peruvian Upwelling Ecosystem (Oliveros-Ramos et al. 2015). The OSMOSE model is a multispecies spatially-explicit individual-based model implemented in the Java language. The modelled area ranges from 20°S to 6°N and 93°W to 70°W, with 1/6° of spatial resolution. The model explicits the life history and spatio-temporal dynamics of 9 species (1 macro-zooplankton group, 1 crustacean, 1 cephalopod and 6 fish species), between 1992 and 2008. The data used to calibrate the model were time series of abundance indices, landings and catch-at-length data. The objective function used a penalized likelihood approach, combining log-normal and multinomial likelihoods. A total of 307 parameters

were estimated in 4 sequential phases. Each calibration trial lasted 5 days using a High Performance Computing (HPC) cluster (64 cores) under the Portable Batch System (PBS) for jobs scheduling. The scripts for this calibration which may be adapted for applications of the `calibrar` package in HPC systems are available at <https://github.com/osmose-model/calibrar.osmose>.

Comparison with other softwares

Implementation of general purpose optimizers can be found in R (see Optimization and Mathematical Programming Task View at CRAN: <http://cran.r-project.org/web/views/Optimization.html>). Two very useful features for model calibration are the performance of constrained optimization (limiting the search to a box by defining lower and upper boundaries to parameter values) and the calibration in multiple phases (to improve the search of the global minimum by performing a sequential approximation). The former option is implemented in several R packages, including the `optim` function (providing the "L-BFGS-B" method, Byrd et al. 1995) and several others wrapped in the `optimx` package (Nash and Varadhan 2011). The latter option is available in some R packages (e.g. `Rcgmin`, `Rvmmin` and `bnmle`) for a single optimization, but a sequential calibration, as described here, would have to be performed manually. For the particular purpose of the calibration of stochastic models, several meta-heuristic and non-derivative based algorithms are now available in R, from EAs (e.g. `genalg`, `DEoptim` and `cmaes` packages) to other nature-inspired algorithms (e.g. Simulated Annealing 'SANN' method in `optim` and the Particle Swarm Optimization (PSO) algorithm in the `hydroPSO` package, Zambrano-Bigiarini and Rojas 2013). However, while all of these packages and algorithms provide support for constrained optimization, none of them provides support for keeping fixed parameters during the course of a single optimization, and multiple phase calibration would have to be performed manually by modifying the objective function for each trial. Furthermore, from the implementation point of view, a very important feature for the calibration of complex models is the parallel implementation of the optimization routine. The PSO algorithm in the `hydroPSO` has its parallel implementation tied to the core of the function and does not allow its use in high-performance clusters, especially under a queue system (e.g. TORQUE), and only the `DEoptim` package provides a more flexible externally configured parallelization.

Additionally, in the construction of the objective function, `calibrar` allows an easy transferability of the calibration problem to other general purpose optimizers, which can be useful under certain circumstances (e.g. see Bolker et al. 2013). There is indeed "no free lunch" in optimization, and no optimization algorithm will perform better than all others for every type of optimization problems (Wolpert and Macready 1997). Other calibration oriented packages like `hydroPSO` provide functions to write parameters and read outputs, but this approach breaks the "objective function" approach for the optimization, and while the `hydromad` package (Andrews et al. 2011) offers support for the automated construction of an objective function in a standard way, it is restricted to some particular cases useful in hydrological modelling. In these regards, our calibration package `calibrar` is meant to be generic enough to be used in a variety of optimization problems, including the calibration of complex (i.e. non-linear, with a lot of parameters to estimate) and stochastic models. Three features of `calibrar` render it particularly useful for the calibration of computationally intensive stochastic models: the parallelization of the simulations, the ability to handle replicate simulations in the evaluation of the objective function and the 'restart' option, which allows the calibration of complex models to be handled under restricted access to high performance resources (e.g. clusters with queue systems and fixed walltime).

Conclusions and perspectives

A successful model calibration implies several computational, theoretical and practical challenges. The `calibrar` package intends to provide a framework to simplify the calibration of complex models, in particular stochastic ones, for which fewer developments have been done compared to those for deterministic and differentiable models. We adopted a "black-box" and "non-intrusive" approach, since most complex models

are computationally intensive and most likely implemented in fast low-level languages; their recoding for calibration purposes is not the best option. In the future, more tests with other models and real-world calibration or optimization problems will help to improve the development of the package, its flexibility and the robustness of the optimization algorithm.

Acknowledgements

The authors were partly funded through the French project EMIBIOS (FRB, contract no. APP-SCEN-2010-II). ROR was supported by an individual doctoral research grant (BSTD) from the "Support and training of scientific communities of the South" Department of IRD, managed by Campus France. We thank David Kaplan, Coleen Moloney, and Arnaud Bertrand for their support and advice during the development of this research. We are also grateful to Philippe Verley, Morgane Travers, Michael Hurtado, Laure Velez, Criscely Luján, Sybille Dueri, Arnaud Grüss, and Tosca Ballerini to have tested and helped to improve *calibrar*. All the calibration experiments were performed with the IFREMER-CAPARMOR HPC 600 facilities. This work is a contribution from the Cooperation agreement between the Instituto del Mar del Peru (IMARPE), the University of Cape Town (UCT) and the Institut de Recherche pour le Développement (IRD), through the LMIs DISCOH and ICEMASA.

References

- Andrews F.T., Croke B.F.W. & Jakeman A.J. (2011) An open software environment for hydrological model assessment and development. *Environmental Modelling & Software* 26:1171-1185.
- Bäck T. & Schewefel H.-P. (1993) An Overview of Evolutionary Algorithms for parameter optimization. *Evolutionary Computation* 1:1-23.
- Bartell S.M. (2003) Effective use of ecological modeling in management: The toolkit concept. *Ecological modeling for resource management* (ed V. Dale), pp. 211-220. Springer Verlag, Berlin.
- Black A.J. & McKane, A.J. (2012) Stochastic formulation of ecological models and their applications. *Trend in Ecology and Evolution* 27(6):337-345. DOI: <http://dx.doi.org/10.1016/j.tree.2012.01.014>
- Beyer H.-G. & Schwefel H.-P. (2002) Evolution strategies: a comprehensive introduction, *Natural Computing* 1:3-52.
- Bolker, B.M., Gardner, B., Maunder, M., Berg, C.W., Brooks, M., Comita, L., Crone, E., Cubaynes, S., Davies, T., de Valpine, P., Ford, J., Gimenez, O., Kéry, M., Kim, E.J., Lennert-Cody, C., Magnusson, A., Martell, S., Nash, J., Nielsen, A., Regetz, J., Skaug, H. & Zipkin, E. (2013) Strategies for fitting nonlinear ecological models in R, AD Model Builder, and BUGS. *Methods in Ecology and Evolution* 4: 501–512.
- Byrd, R. H., Lu, P., Nocedal, J. & Zhu, C. (1995) A limited memory algorithm for bound constrained optimization. *SIAM J. Scientific Computing*, 16, 1190–1208.
- Cropper, W.P. Jr. & Anderson, P.J. (2004) Population dynamics of a tropical palm: use of a genetic algorithm for inverse parameter estimation. *Ecological Modelling* 177: 119–127.
- Duboz, R., Versmisse, D., Travers, M., Ramat, E. & Shin, Y.-J. (2010) Application of an evolutionary algorithm to the inverse parameter estimation of an individual-based model. *Ecological Modelling* 221(5):840-849.
- Fournier, D.A., Skaug, H.J., Ancheta, J., Ianelli, J., Magnusson, A., Maunder, M.N., Nielsen, A. & Sibert, J. (2012) AD Model Builder: using automatic differentiation for statistical inference of highly parameterized complex nonlinear models. *Optimization Methods and Software*, 27:2, 233-249, DOI: 10.1080/10556788.2011.597854.
- Griewank, A. & Corliss, G.F. (1992) *Automatic Differentiation of Algorithms: Theory, Implementation, and Application*. SIAM, Philadelphia, PA, USA.
- Hansen, N., & Ostermeier, A. (2001) Completely de-randomized self-adaptation in evolution strategies. *Evolutionary Computation* 9(2):159–195.
- Ionides, E. L., Breto, C. & King, A.A. (2006) Inference for nonlinear dynamical systems. *PNAS* 103(49): 18438–18443.
- Jones, D.R., Schonlau, M. & Welch, W.J. (1998) Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization* 13: 455–492.

- Jones, G. (1998) Genetic and evolutionary algorithms. In Encyclopedia of Computational Chemistry (Schleyer, P. v.R. et al., eds), pp. 1127–1136, John Wiley & Sons.
- Nash, J.C. & Varadhan, R. (2011) Unifying Optimization Algorithms to Aid Software System Users: optimx for R. *Journal of Statistical Software*, 43(9), 1–14. <http://www.jstatsoft.org/v43/i09/>.
- Nash, J.C. & Walker-Smith, M. (1987) *Nonlinear Parameter Estimation: an Integrated System in BASIC*. Marcel Dekker, New York. 493pp.
- Newman, K.B., Fernández, C., Thomas, L. & Buckland, S.T. (2009) Monte Carlo Inference for State–Space Models of Wild Animal Populations. *Biometrics* 65, 572–583.
- Oliveros-Ramos, R., Verley, P. & Shin, Y.-J., 2015. A sequential approach to calibrate ecosystem models with multiple time series data. [arXiv:1509.06123v1](https://arxiv.org/abs/1509.06123v1).
- Poovathingal, S.K. & Gunawan, R. (2010) Global parameter estimation methods for stochastic biochemical systems, *BMC Bioinformatics* 11:414.
- R Core Team (2016) R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <http://www.R-project.org/>.
- Revolution Analytics & Weston, S. (2014) foreach: Foreach looping construct for R. R package version 1.4.2. <http://CRAN.R-project.org/package=foreach>. [accessed 09 June 2014]
- Rios, L.M. & Sahinidis, N.V. (2013) Derivative-free optimization: a review of algorithms and comparison of software implementations. *Journal of Global Optimization* 56:1247–1293.
- Ross, J.V., Pagendam, D.E. & Pollet P.K. (2009) On parameter estimation in population models II: Multi-dimensional processes and transient dynamics. *Theoretical Population Biology* 75: 123–132.
- Shin, Y.-J. & Cury, P. (2001) Exploring fish community dynamics through size-dependent trophic interactions using a spatialized individual-based model. *Aquatic Living Resources*, 14(2): 65–80.
- Shin, Y.-J. & Cury, P. (2004) Using an individual-based model of fish assemblages to study the response of size spectra to changes in fishing. *Canadian Journal of Fisheries and Aquatic Sciences*, 61: 414–431.
- Tashkova, K., Silc, J., Atanasova, N. & Dzeroski, S. (2012) Parameter estimation in a nonlinear dynamic model of an aquatic ecosystem with meta-heuristic optimization. *Ecological Modelling* 226: 36– 61.
- Thiele, J.C., Kurth, W. & Grimm, V. (2012) RNetLogo: an R package for running and exploring individual-based models implemented in NetLogo. *Methods in Ecology and Evolution* 3:480–483.
- Trautmann H., Mersmann O. and Arnu D. (2011) cmaes: Covariance Matrix Adapting Evolutionary Strategy. R package version 1.0-11. <https://CRAN.R-project.org/package=cmaes>
- Travers-Trolet, M., Shin, Y.-J. & Field, J.G. (2013) An end-to-end coupled model ROMS-N2P2Z2D2-OSMOSE of the southern Benguela foodweb: parameterisation, calibration and pattern-oriented validation, *African Journal of Marine Science*, 36:1, 11–29, DOI:10.2989/1814232X.2014.883326
- Walker, D.M., Pérez-Barbería, F.J. & Marion, G. (2006) Stochastic modelling of ecological processes using hybrid Gibbs samplers. *Ecological Modelling* 198:40–52.
- Walter, E. and Pronzato, L. (1997) *Identification of parametric models from Experimental data*. Springer Masson. 413pp.
- Wolpert D.H. & Macready W.G. (1997) No Free Lunch Theorems for Optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82.
- Zambrano-Bigiarini, M. & Rojas, R. (2013) A model-independent Particle Swarm Optimisation software for model calibration. *Environmental Modelling & Software* 43:5–25.

Appendix S1. Full description of the algorithm.

Appendix S2. Scripts with examples.

Supplementary material 1:

An evolutionary strategy for the calibration of ecological models using maximum likelihood estimation

calibrar: an R package for fitting complex ecological models

1 Introduction

Evolutionary algorithms (EA) are computer programs designed for the automatic solving of complex problems such as minimization of functions, and are inspired by the process of Darwinian evolution (Jones 1998). The three main types of EA are Genetic Algorithms (GA), Evolutionary Programming (EP) and Evolutionary Strategies (ES). Historically, Evolutionary Programming and especially Genetic Algorithms were designed with a broader range of applications (Bäck and Schwefel, 1993) while Evolutionary Strategies (ESs) were specifically designed for parameter optimization problems (Jones, 1998). For optimization problems, EAs work over a population of “individuals” searching over the space of solutions. Each individual encodes a solution (e.g. a vector of parameter values for a model) to the problem which performance can be assessed. EAs rely on three main processes: selection, recombination and mutation. The selection process is intended to select the individuals which will produce offspring (i.e the population for the next generation). The recombination process allows inbreeding the selected individuals (parents) in an attempt to enhance their performance. Finally, the mutation process produces random variability in the population, normally by modifying the solution encoded by the parents.

The *calibrar* package includes a novel algorithm based on Evolutionary Strategies (ES), the AHR-ES. The main novelty of the algorithm developed is the implementation of a recombination process that takes into account: i) the variability in the parameters, which provides a better fit for each data type, and ii) the probabilistic nature of the likelihood approach to weight the potential candidates to parameter solutions. Also, a similar approach for self-adaptation as in Hansen and Ostermeier (2001) has been implemented to avoid a premature convergence. These modifications have shown a great increase in performance compared to other ESs used for the calibration of complex stochastic models, like in Duboz et al. (2010). The full technical details of the implementation of the algorithm are described in the next sections.

2 Evolutionary strategies

In ESs, selection and recombination are deterministic parametric processes. Additionally, EAs include some meta-parameters controlling the behavior of the algorithm itself (e.g. the mutation rates). ESs also include “self-adaptation” procedures allowing to make the meta-parameters of the algorithm vary to improve their performance over the evolutionary process. ESs focus on mutation as the main search operator, and it has been pointed out that it is necessary to use recombination in connection to self-adaptation to improve the performance of the algorithm (Thomas and Schwefel, 1993). A comprehensive synthesis of Evolutionary Strategies can be found in Beyer and Schwefel

(2002).

We consider a population $\{x_i\}$, with $x_i \in \mathbb{R}^n$, for $i = 1, \dots, \lambda$ and n the dimension of the problem (i.e the number of parameters to estimate). We also need to define an objective function f (so called fitness function) to be minimized. So, for each x_i we can calculate $f(x_i)$ and we can sort the individuals of the population by their fitness values:

$$f(x_{1:\lambda}) \leq f(x_{2:\lambda}) \leq \dots \leq f(x_{\lambda:\lambda}) \quad (1)$$

Where $x_{i:\lambda}$ encodes the i -th lower value for the function f among the population. This allows us to carry on the selection of the best $\mu < \lambda$ individuals of the population, which will constitute the parents for the next generation.

The recombination of the parents can follow different rules. It can be as simple as taking the mean (or weighted mean) of the μ selected parents. Finally, the mutation process is used to produce a new generation, for example by sampling the new x_i from a multivariate normal distribution:

$$x_i \sim N(m, C)$$

where m is an n -dimensional vector resulting from the recombination of the parents and C is a covariance matrix. During the course of the evolutionary process, m will converge to an optimal solution.

In the algorithm developed in this work, we introduce a new method for an adaptative hierarchical recombination (AHR), optimized for parameter estimation of models using several sources of information (i.e calibration using several sources of data). Additionally, in order to improve the convergence and search capabilities, we implement self-adaptation procedures to improve the adaptation of the covariance matrix C during the optimization.

In order to introduce a self-adaptation process in our algorithm, we assume C is a diagonal matrix, while extending the results to a generic covariance matrix is a work under progress. In the next section, the algorithm developed is described in detail.

3 The AHR-ES Algorithm

3.1 Objective function

We are considering a general class of objective functions f :

$$f(x) = f_0(x) + \sum_{k=1}^K f_k(x), \quad (2)$$

where $x \in \mathbb{R}^n$ is a parameter vector and f_k , $k = 0, \dots, K$ are the *partial fitnesses*. The objective of the calibration is to optimize $f(x)$, the search being directed by the recombination between individuals with “local” success (optimizing f_k , $k = 1, \dots, K$).

It is important to notice that we are not sorting the parents according to the partial fitness for the f_0 component, but this component contributes to the total fitness and the initial selection. In particular, f_k could be the likelihood function associated to each calibration variable. By using likelihood

functions, it is straightforward to build fitness functions to calibrate variables with data time series. Also, this choice makes a handful of statistical procedures available to test the goodness of fit, to estimate confidence intervals, etc. On the other hand, likelihood fitness functions could be very complex and highly multimodal, especially when handling a model with non-linear relationships and stochasticity. Optimizing likelihood functions for complex models could be prone to premature stagnation and requires more generations to find optimal solutions, reason why it is important to reduce population sizes (to reduce computing time) and to use properly defined self-adapted mutation rates (to increase rate of convergence).

3.2 Selection

We select the $\mu < \lambda$ parents \tilde{x}_i ($i = 1, \dots, \mu$) which have the lowest value of the objective function f . Then, for each partial fitness f_i ($i = 1, \dots, K$) we will sort the parents as in Equation (1):

$$f_k(\tilde{x}_{1;\mu,k}) \leq f_k(\tilde{x}_{2;\mu,k}) \leq \dots \leq f_k(\tilde{x}_{\mu;\mu,k}) \quad (3)$$

for each $m = 1, \dots, K$ partial fitness.

3.3 Recombination

As a first step, we will recombine the parents according to their success at optimizing each partial fitness f_k , given a set of weights w_i ($i = 1, \dots, \mu$):

$$x_k = \sum_i w_i \tilde{x}_{i;\mu,m} \quad (4)$$

$$\sigma_k^2 = \sum_i w_i \tilde{x}_{i;\mu,m}^2 - x_k^2 \quad (5)$$

such that $w_i \geq w_j$ for $i < j$ and $\sum_i w_i = 1$. Note that x_k^2 is taken entry-wise, i.e. squaring each component of x_k independently (Hadamard product). This initial recombination allows to better use the information in all selected individuals, and particularly to reduce the impact of selecting an individual with a good fitness value just by chance, especially when dealing with stochastic models. As part of the recombination we also calculate σ_k which provides information about the variability of each parameter value among the parents.

Then, we exploit all the historical information on x_k and σ_k by exponentially weighting the past of the recombined parents:

$$\mathbf{x}_k(g) = (1 - \alpha) \mathbf{x}_k(g-1) + \alpha x_k \quad (6)$$

$$\mathbf{s}_k^2(g) = (1 - \alpha) (\mathbf{x}_k^2(g-1) + \mathbf{s}_k^2(g-1)) + \alpha (\sigma_k^2 + x_k^2) - \mathbf{x}_k^2(g) \quad (7)$$

for each $m = 1, \dots, K$ partial fitness, and generation g . Here, $\mathbf{x}_k(g)$ and $\mathbf{s}_k(g)$ are calculated as moving average and variance for generation g , to take into account past information with exponentially decreasing weights given by $\alpha \in [0, 1]$, a meta-parameter of the algorithm, which controls the rate at which the algorithm learns from the current parents. Particularly, \mathbf{s}_k provides information on how important a particular parameter is for the minimization of the objective function, since the

more important the parameter the smaller the variability that we would expect across the generations. Now, let's define $\mathbf{s}_k^{(\min)} = \min_n \mathbf{s}_k$ and $\mathbf{s}_k^{(\max)} = \max_n \mathbf{s}_k$, the minimum and maximum over the n entries of \mathbf{s}_k , respectively to calculate:

$$\hat{w}_k = \left[\frac{\mathbf{s}_k^{(\max)} - \mathbf{s}_k}{\mathbf{s}_k^{(\max)} - \mathbf{s}_k^{(\min)}} \right]^\beta \quad (8)$$

$$w_k = \frac{\hat{w}_k}{\|\hat{w}_k\|_1} \quad (9)$$

for $\beta \geq 1$, and $\|\hat{w}_k\|_1$ is the L_1 norm of \hat{w}_k , taken to make the sum of w_k equal to 1. Again, the quotient and the power are taken entry-wise. w_k ponderates the relative importance of each parameter to the partial fitness m . When parameters are bounded, the vector \mathbf{s}_k can be divided by the ranges of each parameter before the recombination stage for rescaling purposes.

Finally, we recombine all parents to produce the *parental genotype* \mathbf{x} by using the weights given by w_k and the first recombined parents given by \mathbf{x}_k :

$$\mathbf{x}[i] = \frac{\sum_{m=1}^M w_k[i] \mathbf{x}_k[i]}{\sum_{m=1}^M w_k[i]}, \quad (10)$$

where $i = 1, \dots, n$ represents the position of a particular parameter in the vectors.

This final recombination uses dynamically changing weights which take into account the variability of each parameter independently and its importance to minimize every partial component of the objective function.

3.4 Mutation

The new individuals of the population in generation $g+1$ will be produced by mutating the parental genotype \mathbf{x} using a multivariate normal distribution:

$$\mathbf{x}_i^{(g+1)} \sim N(\mathbf{x}^{(g)}, \sigma_{size}^{(g)} C^{(g)}) \quad (11)$$

for $i = 1, \dots, \lambda$. The matrix $C^{(g)}$ is constructed following the self-adaptation algorithm techniques (Covariance Matrix Adaptation CMA-ES; Hansen and Ostermeier 2001) and σ_{size} is the step size control calculated as in Hansen and Ostermeier (2001). The reader can read the source code for details on this particular implementation.

Additionally, when the parameters are bounded, a truncated multivariate normal distribution is used for the mutation process instead of a multivariate normal distribution.

Supplementary material 2:

Scripts with examples

calibrar: an R package for fitting complex ecological models

1 Implementation of a simple example

To illustrate the use of this function, we will try to minimize the Sphere function (with random noise), defined as:

$$F(x) = \sum_{i=1}^n x_i^2 + e_i,$$

where $x = (x_1, x_2, \dots, x_n)$ and $e_i \sim N(0, \sigma)$. This function has a minimum expected value of 0 at the origin. The two obligatory arguments of the calibrate function, with no default values, are `par` and `fn`, i.e. the starting parameter value for the search and the function to minimize, respectively (see Table 2 in the main text or `?calibrate`). For $n = 5$, the minimization can be run as follows:

```
calibrate(par=rep(NA, 5), fn=Sphere)
```

When NA (not available) values are provided as initial search points, the function will try to choose an appropriate start value (see the help page of the function for details). However, for a real calibration problem, providing a good start value (based on prior knowledge of the parameters) would improve the performance of the calibration (Bolker et al. 2013), even when using a global optimization algorithm as in `calibrar`.

As the objective function is stochastic, the search surface depends on the particular realization of the random variables involved. Here we can specify the number of replicate simulations we want to run for a particular set of parameters, and the expected value of the objective function will be the average over the replicates.

```
calibrate(par=rep(NA, 5), fn=Sphere, replicates=3)
```

It is possible to provide, additionally, lower and upper bounds for the parameters

```
calibrate(par=rep(0.5, 5), fn=Sphere, replicates=3, lower=-5, upper=5)
```

If only one value is provided for lower and upper instead of a vector, it will be used for all parameters (with a warning). Finally, the phases argument indicates whether the calibration is run in multiple phases, by specifying at which phase the parameters are included in the optimization:

```
calibrate(par=rep(0.5, 5), fn=Sphere, replicates=3, lower=-5, upper=5,  
          phases=c(1,1,1,2,3))
```

This call will perform three sequential optimizations. In the first one, only the first three parameters are estimated, so the last ones remain constant at the start value (0.5). In the second phase, the fourth parameter becomes activated, and a second optimization is carried out for estimating the first four parameters, and keeping the last one constant. The main difference from the first phase is that the starting points for the first three parameters are not from the original set of starting values defined by `par`, but are the optimal values obtained from the first optimization phase. Lastly, a

third and final optimization is carried out with all the parameters, starting from the optimal values obtained in the second phase. Negative integers and NA are accepted for phases, both meaning that the parameter will never be active, so it will remain constant throughout the calibration. This can be particularly useful for tests with simpler models where some parameters remain constant, without needing to change the objective function. Additionally, a different number of replicates can be indicated for each phase. Since the objective of the initial phases is to try to get an improved vector of start values for a final calibration with all the parameters, it can be useful to reduce the computer time by using fewer replicates in the initial phases.

```
calibrate(par=rep(0.5, 5), fn=Sphere, replicates=c(1,1,4), lower=-5, upper=5,
          phases=c(1,1,1,2,3))
```

The default value for the replicates is 1, since EAs can handle the optimization of stochastic functions directly, but by using more replicates in the last phase we reduce the stochasticity of the surface, which can help to estimate the optimal parameters for very noisy problems.

In the next sections we provide some scripts useful to test the main functionalities of the package.

2 Parallel execution and restart functionality

```
# Restarting a calibration -----

# this calibration save results on the disk for restart purposes
calibrate(par=rep(0.5, 5), fn=SphereN, replicates=3, lower=-5, upper=5,
          phases=c(1,1,1,2,3), control=list(restart.file="sphere"))
# this calibration take no time, because starts from (already finished)
# previous one
calibrate(par=rep(0.5, 5), fn=SphereN, replicates=3, lower=-5, upper=5,
          phases=c(1,1,1,2,3), control=list(restart.file="sphere"))

# Parallel execution -----

nCores = 6 # number of cores to be used
myCluster = makeCluster(nCores)
registerDoSNOW(myCluster) # register the parallel backend
# this is slower than sequential for very fast models
calib = calibrate(par=rep(0.5, 5), fn=SphereN,
                  replicates=3, lower=-5, upper=5,
                  phases=c(1,1,1,2,3),
                  control=list(parallel=TRUE, nCores=nCores))
stopCluster(myCluster) # close the parallel connections
```

3 A simple linear model fitting as benchmarking

```
require(calibrar)
```

```
N = 9 # number of variables in the linear model
```

```

T = 100 # number of observations
noise = FALSE # add gaussian noise to the model
shift = FALSE # add a random shift to the slopes
sd = runif(1) # standard deviation of the gaussian noise

# observed data
x = t(matrix(rnorm(N*T, sd=sd), nrow=N, ncol=T))

# slopes for the linear model (real parameters)
slope = seq_len(N) + shift*sample(c(-10, 10), N, replace=TRUE)
# intercept for the linear model (real parameters)
intercept = pi
# real parameters
real = list(intercept=intercept, slope=slope)

# function to simulate the linear model
linear = function(x, par) {
  stopifnot(length(x)==length(par$slope))
  out = sum(x*par$slope) + par$intercept
  return(out)
}

# simulated data
y = apply(x, 1, linear, par=real) + noise*rnorm(nrow(x), sd=mean(sd))

# objective function (residual squares sum)
obj = function(par, x, y) {
  y_sim = apply(x, 1, linear, par=par)
  out = sum((y_sim - y)^2)
  return(out)
}

lower = relist(rep(-10, N+1), skeleton=start)
upper = relist(rep(+10, N+1), skeleton=start)

# initial guess for optimization
start = list(intercept=0, slope=rep(0, N))

cat("Running optimization algorithms\n")
cat("\t", date(), "\n")

cat("Running calibrar AHR-ES (unconstrained)\n")
print(system.time(es <- calibrate(par=start, fn=obj, x=x, y=y)))

cat("Running calibrar AHR-ES (constrained)\n")
print(system.time(es2 <- calibrate(par=start, fn=obj, x=x, y=y,
                                   lower=lower, upper=upper)))

cat("Running linear model\n")
print(system.time(mod <- lm(y ~ x)))

```

```

cat("Running optim CG\n")
print(system.time(cg <- calibrate(par=start, fn=obj, x=x, y=y, method="CG")))

cat("Running optim SANN\n")
print(system.time(sann <- calibrate(par=start, fn=obj, x=x, y=y,
method="SANN")))

cat("Running optimx Nelder-Mead\n")
print(system.time(nm <- calibrate(par=start, fn=obj, x=x, y=y,
method="Nelder-Mead")))

cat("Running optimx BFGS\n")
print(system.time(bfgs <- calibrate(par=start, fn=obj, x=x, y=y,
method="BFGS")))

cat("Running cmaes CMA-ES\n")
print(system.time(cma <- calibrate(par=start, fn=obj, x=x, y=y,
lower=lower, upper=upper, method="cmaes")))

final = rbind(real=unlist(real),
'AHR-ES' = unlist(es$par),
'AHR-ES (constrained)' = unlist(es2$par),
lm=coef(mod),
SANN = unlist(sann$par),
'CMA-ES' = unlist(cma$par),
'Nelder-Mead' = unlist(nm$par),
'BFGS' = unlist(bfgs$par),
CG = unlist(cg$par))

print(final)

```

4 Fitting an Predator-Prey model

```

require(calibrar)
set.seed(880820)
path = NULL # NULL to use the current directory
# create the demonstration files
demo = calibrarDemo(path=path, model="PredatorPrey", T=100)
# get calibration information
calibrationInfo = getCalibrationInfo(path=demo$path)
# get observed data
observed = getObservedData(info=calibrationInfo, path=demo$path)
# Defining 'runModel' function
runModel = calibrar:::PredatorPreyModel
# real parameters
cat("Real parameters used to simulate data\n")
print(unlist(demo$par)) # parameters are in a list
# objective functions
obj = createObjectiveFunction(runModel=runModel, info=calibrationInfo,
observed=observed, T=demo$T)

```

```

obj2 = createObjectiveFunction(runModel=runModel, info=calibrationInfo,
    observed=observed, T=demo$T, aggregate=TRUE)
cat("Starting calibration...\n")
cat("Running optimization algorithms\n", "\t")
cat("Running optim AHR-ES\n")
ahr = calibrate(par=demo$guess, fn=obj, lower=demo$lower, upper=demo$upper,
    phases=demo$phase)
cat("Running optim CG\n")
cg = calibrate(par=demo$guess, fn=obj2, phases=demo$phase, method="CG")
cat("Running optimx BFGS\n")
lbfgsb = calibrate(par=demo$guess, fn=obj2, lower=demo$lower,
    upper=demo$upper, phases=demo$phase, method="L-BFGS-B")
cat("Running cmaes CMA-ES\n")
cma = calibrate(par=demo$guess, fn=obj2, lower=demo$lower, upper=demo$upper,
    phases=demo$phase, method="cmaes")
cat("Running optim SANN\n")
nm = calibrate(par=demo$guess, fn=obj2, phases=demo$phase,
    method="Nelder-Mead")

comps = summary(demo, ahr, cg, lbfgsb, cma, nm)
print(comps)

```

5 Fitting an Autoregressive Poisson mixed model

```

require(calibrar)
set.seed(880820)
path = NULL # NULL to use the current directory
# create the demonstration files
demo = calibrarDemo(path=path, model="PoissonMixedModel", L=5, T=25)
# get calibration information
calibrationInfo = getCalibrationInfo(path=demo$path)
# get observed data
observed = getObservedData(info=calibrationInfo, path=demo$path)
# read forcings for the model
forcing = read.csv(file.path(demo$path, "master", "environment.csv"),
    row.names=1)
# Defining 'runModel' function
runModel = function(par, forcing) {
    output = calibrar:::.PoissonMixedModel(par=par, forcing=forcing)
    output = c(output, list(gammas=par$gamma)) # adding gamma parameters for
    penalties
    return(output)
}
# real parameters
cat("Real parameters used to simulate data\n")
print(demo$par)
# objective functions
obj = createObjectiveFunction(runModel=runModel, info=calibrationInfo,
    observed=observed, forcing=forcing)
obj2 = createObjectiveFunction(runModel=runModel, info=calibrationInfo,
    observed=observed, forcing=forcing, aggregate=TRUE)

```

```

cat("Starting calibration...\n")
control = list(weights=calibrationInfo$weights, maxit=2e5) # control
parameters
cat("Running optimization algorithms\n", "\t", date(), "\n")
cat("Running optim AHR-ES\n")
ahr = calibrate(par=demo$guess, fn=obj, lower=demo$lower, upper=demo$upper,
               control=control)
cat("Running optim CG\n")
cg = calibrate(par=demo$guess, fn=obj2, method="CG", control=control)
cat("Running optimx BFGS\n")
lbfgsb = calibrate(par=demo$guess, fn=obj2, lower=demo$lower,
                  upper=demo$upper, method="L-BFGS-B", control=control)
cat("Running cmaes CMA-ES\n")
cma = calibrate(par=demo$guess, fn=obj2, lower=demo$lower, upper=demo$upper,
               method="cmaes", control=control)
cat("Running optim SANN\n")
sann = calibrate(par=demo$guess, fn=obj2, method="SANN", control=control)

comps = summary(demo, ahr, cg, lbfgsb, cma, sann)
print(comps)

```